IPv6 Security Issues

Samuel Sotillo
East Carolina University
ss0526@ecu.edu

Abstract

Deployment of a new generation of Internet protocols is on its way. It is a process that may take several years to complete. In the meantime, the deployment raises considerable new issues, being security one of the most compelling.

This paper reviews some of the improvements associated with the new Internet Protocol version 6, with an emphasis on its security-related functionality. At the end, it concludes summarizing some of the most common security concerns the new suite of protocols creates.

1. Introduction

For more than three decades now, the Internet's end-to-end model has functioned remarkably well. This model has allowed the evolution of a transparent network architecture that efficiently supports the transport of data without caring what the data it-self represents [7]. Furthermore, being transparent and application-neutral has facilitated the creation and evolution of new Internet applications and services that operate on the same thirty-something network architecture—which until recently had not required any major overhaul.

Unfortunately, the landscape is changing. Today, the Internet has grown to be a million-network network, which is something with startling consequences. For instance, one of the most publicized consequences of this growth has been the depletion of the Internet's address space. Initially, the Internet's address space consisted of 2^{32} addresses—about 4 billion addresses. Today, however, that amount is insufficient, even more if we consider emerging new technologies such as 3G/4G wireless devices and other wireless appliances [1].

Another consequence of the current Internet's exponential growth has to do with security [3]. At the time of its design, and keeping up with the original end-to-end model, the Internet was thought as a "friendly" environment. Therefore, no security was embedded in the original architecture [4]. Today, the new million-network network has become a very "hostile" environment. Although important new techniques have been introduced to overcome some of the Internet's best known security deficiencies (SSL, IPSec, etc.), they seem to be insufficient. Unfortunately, despite all recent improvements, the underlying infrastructure of the Internet continues to lack the appropriate security framework.

Aware of the limitations of the current Internet infrastructure, which is based on the Internet Protocol version 4 (IPv4) suite of protocols, the Network Working Group of the Internet Engineering Task Force (IETF) proposed in 1998 a new suite of protocols called the Internet Protocol version 6 (IPv6) [9]. This new suite of protocols addresses several of the issues that affect IPv4-based networks, including its lack of network level security. In this paper, we outline the advantage that, in the matter of security, the new protocol suite brings to the table. Also, we review some of the challenges the new protocol suite faces as the Internet continues its quest for global domination.

2. IPv4 security issues

Before studying IPv6, we need to understand some of the best known limitations of its predecessor, IPv4. As mentioned before, IPv4 was designed with no security in mind. Because of its end-to-end model, IPv4 assumes that security should be provided by the end nodes [7]. For instance, if an application such as e-mail requires encryption services, it should be the responsibility of such application at the end nodes to provide such services. Today, the original Internet continues to be completely transparent and no security framework provides for resilient against threats such as:

- *Denial of service attacks (DOS)*: in this kind of attack certain services are flooded with a large amount of illegitimate requests that render the targeted system unreachable by legitimate users. An example of DOS attack that results from an architectural vulnerability of IPv4 is the broadcast flooding attack or Smurf attack [12]
- *Malicious code distribution*: viruses and worms can use compromised hosts to infect remote systems. IPv4's small address space can facilitate malicious code distribution [12].
- *Man-in-the-middle attacks*: IPv4's lack of proper authentication mechanisms may facilitate men-in-

the-middle attacks. Additionally, ARP poisoning (see below) and ICM redirects can also be used to perpetrate this type of attacks [12] [2].

- *Fragmentation attacks*: this type of attacks exploits the way certain operating systems handle large IPv4 packets. An example of this type of attack is the *ping of death* attack. In a *ping of death* attack the target system is flooded with fragmented ICMP *ping* packets. With each fragment, the size of the reassembled *ping* packet grows beyond the packet size limit of IPv4—therefore, crashing the target system [12].
- *Port scanning and other reconnaissance attacks*: in this type of attacks a whole section of a network is scanned to find potential targets with open services. Unfortunately, IPv4's address space is so small that scanning a whole class C network can take a little more than 4 minutes [13].
- *ARP poisoning and ICMP redirect*: in IPv4 networks, the Address Resolution Protocol (ARP) is responsible for mapping a host's IP address with its physical or MAC address. This information is stored by each host in a special memory location known as the ARP table. Each time a connection with an unknown host is needed, an ARP request is sent out on the network. Then, either the unknown host responds broadcasting its own IP address or a router does it with the appropriate information. *ARP poisoning* occurs when forged ARP responses are broadcasted with incorrect mapping information that could force packets to be sent to the wrong destination. A similar approach is used by ICMP redirect attacks [12].

However, many techniques have been developed to overcome some of the IPv4 security limitations. For instance, although Network Address Translation (NAT) and Network Address Port Translation (NAPT) were introduced to facilitate the re-use and preservation of a rapidly depleting IPv4 address space, these techniques can provide also for certain level of protection against some of the aforementioned threats [11]. Also, the introduction of IPSec facilitated the use of encryption communication, although its implementation is optional and continues to be the sole responsibility of the end nodes.

3. IPv6 in a nutshell

First of all, it is important to emphasize that IPv6 is not a superset of IPv4 but an entirely new suite of protocols. For that reason, and because of space limitations, we will only summarize some of IPv6 most interesting features [1] [2]:

- *Larger address space*: as mentioned above, IPv4 provides as many as 2^{32} addresses[1]. On the other hand, IPv6 provides for as many as 2^{128} addresses[2].
- *Hierarchical addressing*: in IPv6 there are three major types of addresses: unicast, multicast, and anycast addresses. Unicast addresses are assigned to a single IPv6 node. Multicast addresses are assigned to multiples nodes within a single multicast group. Packets sent to a multicast address must be delivered to all members of the same multicast group. On the other hand, although anycast addresses are also assigned to groups of nodes, they do not need to be delivered to all members of the group—it is sufficient that one node receives the packets. Additionally, IPv6 defines a new routing infrastructure that provides for more efficient and smaller routing tables.
- *Stateless and stateful address configuration*: IPv6 allows hosts to acquire IP addresses either in a stateless or autonomous way or through a controlled mechanism such as DHCPv6.
- *Quality-of-service*: the IPv6 packet header contains fields that facilitate the support for QoS for both differentiated and integrated services.
- *Better performance*: IPv6 provides for significant improvements such as better handling of packet fragmentation, hierarchical addressing, and provisions for header chaining that reduce routing table size and processing time.
- *Built-in security*: although IPSec is also available for IPv4 implementations, it is not mandated but optional. Support for IPSec in IPv6 implementations is not an option but a requirement.
- *Extensibility*: despite the fact that IPv6 addresses are four times larger than IPv4 addresses, the new IPv6 header is just twice the size of the IPv4 header (i.e., two times 20 bytes = 40 bytes). As shown in Fig 1, the new IPv6 header does not include any optional fields. It does not include a checksum either. Optional fields can be added as extension headers up to the size of the IPv6 packet. This feature does not only provide for better extensibility but also for reducing the time

[1] Actually, the amount of available IPv4 addresses is considerably smaller. Unfortunately, the initial class-based scheme used to assign IP addresses resulted in a very inefficient use of the original 4,294,967,296 addresses, contributing to the ultimate depletion problem we face today [4].

[2] 2^{128} represents about $3.402823669209384634633746074317 7 \times 10^{38}$ unique addresses (theoretically).

a router process IPv6 header options, increasing the network overall performance.

- *Mobility*: IPv6 provides mechanisms that allow mobile nodes to change their locations and addresses without loosing the existing connections through which those nodes are communicating. This service is supported at the Internet level and therefore is fully transparent to upper-layer protocols.

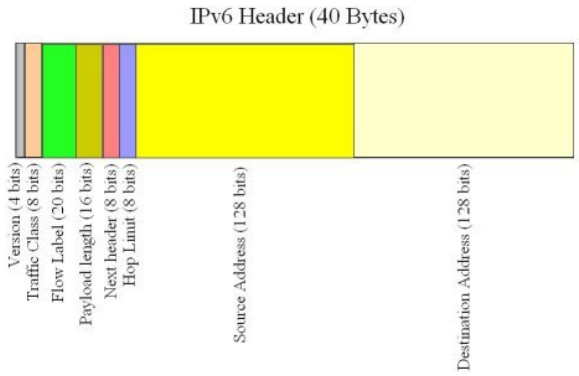

Figure 1: Internet Protocol version 6 header

Of course, IPv6 offers many other interesting features that are beyond the scope of this paper. In the following sections we focus on whether these new set of protocols are better prepared for facing some of today's more pressing security issues on the Internet.

4. IPv6 security improvements

As noticed in [2], it is important to begin this section acknowledging that IPv6 is not necessarily more secure than IPv4. In fact, IPv6 approach to security is only marginally better than IPv4 but not radically new [4]. The following sub-sections summarize some IPv6's improvements that provide for better network security.

4.1. Large address space

Port scanning is one of the best known reconnaissance techniques in use today. Port scanning allows "black-hats" to listen to specific services (ports) that could be associated to well-known vulnerabilities [13].

In IPv4 networks, port scanning is a relatively simple task. Most IPv4 segments are Class C, with 8 bits allocated for host addressing. Scanning a typical IPv4 subnet, at a rate of one host per second, translates into:

$$2^8\, hosts \times \frac{1\text{sec}\,ond}{1 host} \times \frac{1\text{min}\,ute}{60\,\text{sec}\,onds} = 4.267 \text{ minutes}$$

In IPv6 networks, the landscape is radically different. IPv6 subnets use 64 bits for allocating host addresses. Consequently, a typical IPv6 subnet requires:

$$2^{64} hosts \times \frac{1\text{sec.}}{1 host} \times \frac{1 year}{31{,}536{,}000 \text{sec.}} = 584{,}942{,}417{,}355 \text{years}$$

Scanning such a large address space is almost an impossible task [2]. However, it is not absolutely impossible [4].

4.2. IPSec

As mentioned above, IPv4 also offers IPSec support. However, IPv4's support for IPSec is optional. By contrast, the RFC4301 mandates for IPv6 to use IPSec in all nodes [2] [10].

IPSec consists of a set of cryptographic protocols that provide for securing data communication and key exchange. IPSec uses two wire-level protocols, *Authentication Header* (AH) and *Encapsulating Security Payload* (ESP). The first protocol provides for authentication and data integrity. The second protocol provides for authentication, data integrity, and confidentiality [10]. In IPv6 networks both the AH header and the ESP header are defined as extension headers. Additionally, IPSec provides for a third suite of protocols for protocol negotiation and key exchange management known as the *Internet Key Exchange* (IKE). This protocol suite provides the initial functionality needed to establish and negotiating security parameters between endpoints. Additionally, it keeps track of this information to guarantee that communication continues to be secure up to the end.

4.2.1. Authentication Header. As mentioned before, the *authentication header* prevents IP packets from being tampered or altered. In a typical IPv4 packet, the AH is part of the payload. Figure 2 shows an example of an IPv4 packet with an AH as payload [14] [15].

When the AH protocol was implemented, there was some concern about how to integrate it to the new IPv6 packet format. The problem centered on the fact that IPv6 extension headers can change in transit as information they contain is updated through the network. To solve this problem, IPv6 AH was designed with flexibility in mind—the protocol authenticates and do integrity check only on those fields in the IPv6 packet header that do not change in transit. Also, in IPv6 packets, the AH is intelligently located at the end of the header chain—but ahead of any ESP extension header or any higher level

protocol such as TCP/UDP [1]. A typical sequence of IPv6 extension headers is shown in Figure 2.

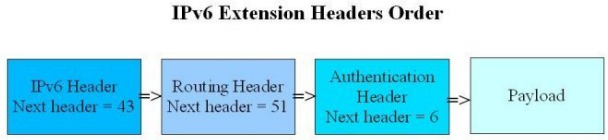

Figure 2: Extension headers order

The AH header protocol also provides optional protection against replay attacks. The protocol uses its sequence number field as part of a sliding window mechanism that prevents arbitrary packet delays and malicious replay [1] [15].

Figure 3 shows a typical AH header.

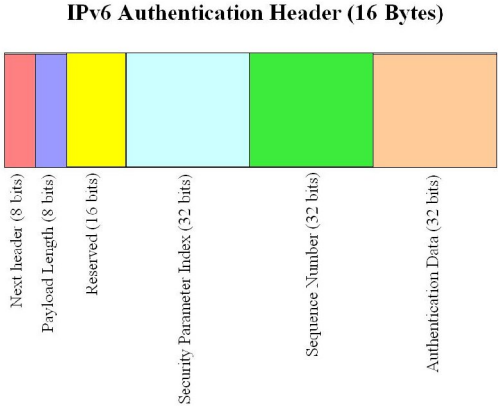

Figure 3: AH header specification

4.2.2. Encapsulating Security Payload. In addition to providing the same functionality the AH protocol provides—authentication, data integrity, and replay protection—ESP also provides confidentiality. In the ESP extension header, the *security parameter index* (SPI) field identifies what group of security parameters the sender is using to secure communication. ESP supports any number of encryption mechanisms. However, the protocol specifies DES-CBC as its default. Also, ESP does not provide the same level of authentication available with AH. While AH authenticates the whole IP header (in fact, only those fields that do not change in transit), ESP authenticates only the information that follows it [1].

ESP provides data integrity by implementing an *integrity check value* (ICV) that is part of the ESP header trailer—the authentication field. The ICV is computed once any encryption is complete and it includes the whole ESP header/trailer—except for the authentication field, of course. The ICV uses *hash message authentication code* (HMAC) with SHA-1 and MD5 as the recommended cryptographic hash functions [15].

Figure 4 shows a typical ESP extension header.

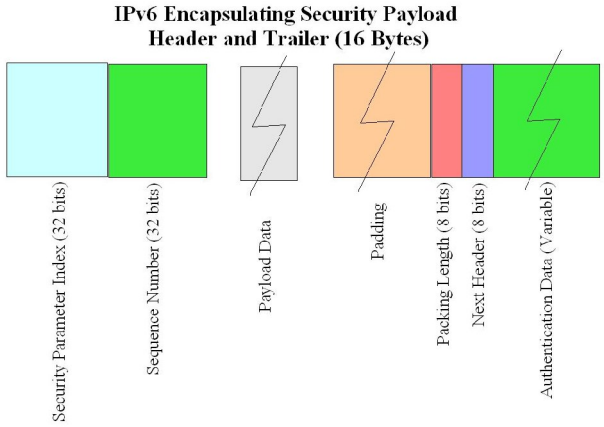

Figure 4: ESP header specification

4.2.3. Transport and tunnel modes. In IPv4 networks, IPSec provides two *modes* of securing traffic. The first one is called *transport* mode and it is intended to provide secure communication between endpoints by securing only the packet's payload. The second one is called *tunnel* mode and it is intended to protect the entire IPv4 packet. However, in IPv6 networks, there is no need for a *tunnel* mode because, as mentioned above, both the AH and ESP protocols provide enough functionality to secure IPv6 traffic [2].

4.2.4. Protocol negotiation and key exchange management. In addition to AH and ESP, IPSec also specifies additional functionality for protocol negotiation and key exchange management [1]. IPSec encryption capabilities depend on the ability to negotiate and exchange encryption keys between parties. To accomplish this task, IPSec specifies an *Internet key exchange* (IKE) protocol. IKE provides the following functionality:

- Negotiating with other people the protocols, encryption algorithms, and keys, to use.
- Exchanging keys easily, including changing them often.
- Keeping track of all these agreements.

To keep track of all protocol and encryption algorithm agreements, IPSec uses the SPI field in both the AH and ESP headers. This field is an arbitrary 32-bit number that represents a *security association* (SA). When communication is negotiated, the receiver node assigns an available SPI which is not in use, and preferably one that has not been used in a while. It then communicates this SPI to its communication partner establishing a *security association*. From then until that SA expires, whenever a node wishes to communicate with the other using the same SA, it must use the same SPI to specify it.

The other node, on receipt, would look at the SPI to determine which SA it needs to use. Then it authenticates and/or decrypts the packet according to the rules of that SA, using the agreed-upon keys and algorithms the SA specifies. The node then uses the same agree-upon information to verify that the data really does come from the node it claims. Also, the node uses the same information to verify that the data has not been modified as well as that no one between the two nodes has read the exchanged data.

Of course, before all this happens, both nodes must negotiate a set of keys. The keys will be used to guarantee that the SA parameters are securely exchanged. IPSec allows for using both automatic and manual key exchange. However, because manual exchange does not scale well, IPSec recommends using IKE. IPSec IKE offers a robust mechanism to authenticate communication parties based on a *public key infrastructure* (PKI). Encryption keys are generated with a Diffie-Hellman algorithm based on each node's public and private key pairs. This mechanism offers perfect forward secrecy (generating keys that are not reliant on previously generated key values) as well as reasonable scalability.

4.3. Neighbor discovery and address auto-configuration

Neighbor discovery (ND) is the mechanism responsible for router and prefix discovery, duplicate address and network unreachability detection, parameter discovery, and link-layer address resolution [1] [4]. This protocol is entirely network-layer based[3]. ND operates in tandem with auto-configuration, which is the mechanism used by IPv6 nodes to acquire either stateful or stateless configuration information. In the stateless mode, all nodes get what they need for global communication, including potential illegal ones. In stateful mode, configuration information can be provided selectively, reducing the possibility for rogue nodes [4]. Both ND and address auto-configuration contribute to make IPv6 more secure than its predecessor. IPv6 provides for TTL values of up to 255; it prevents against outside sourcing of ND packets or duplicate addresses [4].

5. IPv6 security issues

From a security point of view, the new IPv6 protocol stack represents a considerable advance in relation to the old IPv4 stack. However, despite its innumerable virtues, IPv6 still continues to be by far vulnerable. In this section we will review some of the areas of IPv6 where security continues to be an important issue.

[3] IPv4's ARP and RARP are link-layer based protocols.

5.1. Dual-stack related issues

Presently, the Internet continues to be mostly IPv4-based. However, it is reasonable to expect that this scenario will change soon as more and more networks are migrated to the new protocol stack. Unfortunately, migrating millions of networks is going to take quite some time. In the meantime, some form of 6to4 dual-stack will supply the desired functionality [1].

Without a doubt, IPv6-IPv4 dual stacks increase the potential for security vulnerabilities—as a consequence of having two infrastructures with specific security problems. However, most of the issues are not a direct result of specific IPv6 design flaws but mostly a result of inappropriate or careless configuration—see [8] for more details.

5.2. Header manipulation issues

The use of *extension headers* and IPSec can deter some common sources of attack based on header manipulation. However, the fact that EH must be processed by all stacks can be a source of trouble—a long chain of EH or some considerably large-size could be used to overwhelm certain nodes (e.g., firewalls) or masquerade an attack. Best practices recommend to filter out traffic with unsupported services [2].

Spoofing continues to be a possibility in IPv6 networks [4]. However, because of ND, spoofing is only possible by nodes on the same network segment.

The same does not apply to 6to4 transition networks. Although one approach to 6to4 transition is using some form of dual-stack functionality, another approach is using some type of tunneling. Because tunneling requires that a protocol is encapsulated in another, its use could be a source of security problems such as address spoofing—in this case if the spoofed address is used to masquerade an external packet as one that was originated from the inside network [4].

5.3. Flooding issues

Scanning for valid host addresses and services is considerably more difficult in IPv6 networks than it is in IPv4 networks. As mentioned above, to effectively scan a whole IPv6 segment may take up to 580 billion years—because the address space uses 64 bits. However, the larger addressing space does not mean that IPv6 is totally invulnerable to this type of attack. Nor the lack of broadcast addresses makes IPv6 more secure. New features such as multicast addresses continue to be source of problems [5]. Smurf-type attacks are still possible on

multicast traffic. Again, filtering out unnecessary traffic is the recommended best practice [2].

5.3. Mobility

Mobility is a totally new feature of IPv6 that was not available in its predecessor. Mobility is a very complex function that raises a considerable amount of concern when considering security. Mobility uses two types of addresses, the real address and the mobile address. The first is a typical IPv6 address contained in an *extension header*. The second is a temporary address contained in the IP header. Because of the characteristics of this networks (something more complicated if we consider wireless mobility), the temporary component of a mobile node address could be exposed to spoofing attacks on the home agent. Mobility requires special security measures and network administrators must be fully aware of them [2] [4] [6].

5. Conclusion

Without doubt, IPv6 represents a considerable improvement if compared to the old IPv4 protocol stack. The new suite of protocols provides innumerable features that improve both the overall functionality as well as some specific security functions. However, it is far from being a panacea. Although IPv6 offers better security (larger address space and the use of encrypted communication), the protocol also raises new security challenges. Ultimately, the new protocol creates as many new security problems as it solves old ones. And if that is not enough, the transition from the old protocol stack to the new one may present even more challenges, something that will guarantee plenty of fun for security network professionals in the foreseeable future.

7. References

[1] Davies, J., *Understanding IPv6*, Microsoft Press, Redmond, WA, 2003.

[2] Popoviciu C.; Levy-Avegnoli, E.; Grossetete, P., *Deploying IPv6 Networks*, Cisco Press, Indianapolis, IN, 2006.

[3] W. Treese, "The State of Security on the Internet", *Communications of the ACM*, September 2004.

[4] Szigeti, S.; Risztics, P., "Will IPv6 bring better security?," *Proceedings 30th Euromicro Conference, 2004*, vol., 532- 537, 31 Aug.-3 Sept. 2004.

[5] Vives, A.; Palet, J., "IPv6 Distributed Security: Problem Statement," *The 2005 Symposium on Applications and the Internet Workshops, 2005. Saint Workshops 2005*, vol., 18- 21, 31-04 Jan. 2005.

[6] Sierra, J.M.; Ribagorda, A.; Munoz, A.; Jayaram, N., "Security protocols in the Internet new framework," *Proceedings IEEE 33rd Annual 1999 International Carnahan Conference on Security Technology, 1999.* vol., no.pp.311-317, 1999.

[7] Bradner, S., "The End-to-End Security," *IEEE Security & Privacy*, vol., no.pp., 76-79, Mar.-Apr. 2006.

[8] Hiromi, R.; Yoshifuji, H., "Problems on IPv4-IPv6 network transition," *Proceedings of the International Symposium on Applications and the Internet Workshop, Saint 2005*, vol., May 2005.

[9] Deering, S.; Hinden, R., "Internet Protocol Version 6 (IPv6) Specification," *RFC 2460*, Dec. 1998, http://www.ietf.org/rfc/rfc2460.txt.

[10] Kent, S.; Seo, K., "Security Architecture for the Internet Protocol," *RFC 4301*, Dec. 2005, http://tools.ietf.org/html/4301.

[11] Smith, M.; Hunt, R., "Network security using NAT and NAPT," *10th IEEE International Conference on Networks, 2002. ICON 2002*, vol., 355- 360, 2002.

[12] Campbell, P.; Calvert, B.; Boswell, S., *Security+ Guide to Network Security Fundamental*, Thomson, Canada, 2003.

[13] Ford, M., "New Internet Security and Privacy Models Enabled by IPv6," *The 2005 Symposium on Applications and the Internet Workshops, 2005. Saint Workshops 2005*, vol., no.pp. 2-5, 31-04 Jan. 2005.

[14] Karve, A., "IP Security," *IT Architect*, Jan. 1998, http://www.itarchitect.com/shared/article/showArticle.jhtml?articleId=17600993.

[15] Friedl, S., "An illustrated Guide to IPSec," *Unixwiz.net*, Aug. 2005, http://www.unixwiz.net/techtips/iguide-ipsec.html.

8. Copyright Notice